

---

# No-Regret Learnability for Piecewise Linear Losses

---

**Arthur Flajolet**

Operations Research Center  
Massachusetts Institute of Technology  
Cambridge, MA 02139  
flajolet@mit.edu

**Patrick Jaillet**

Dept. of Electrical Engineering and Computer Science  
Operations Research Center  
Massachusetts Institute of Technology  
Cambridge, MA 02139  
jaillet@mit.edu

## Abstract

In the convex optimization approach to online regret minimization, many methods have been developed to guarantee a  $O(\sqrt{T})$  bound on regret for subdifferentiable convex loss functions with bounded subgradients, by using a reduction to linear loss functions. This suggests that linear loss functions tend to be the hardest ones to learn against, regardless of the underlying decision spaces. We investigate this question in a systematic fashion looking at the interplay between the set of possible moves for both the decision maker and the adversarial environment. This allows us to highlight sharp distinctive behaviors about the learnability of piecewise linear loss functions. On the one hand, when the decision set of the decision maker is a polyhedron, we establish  $\Omega(\sqrt{T})$  lower bounds on regret for a large class of piecewise linear loss functions with important applications in online linear optimization, repeated zero-sum Stackelberg games, online prediction with side information, and online two-stage optimization. On the other hand, we exhibit  $o(\sqrt{T})$  learning rates, achieved by the Follow-The-Leader algorithm, in online linear optimization when the boundary of the decision maker's decision set is curved and when 0 does not lie in the convex hull of the environment's decision set. Hence, the curvature of the decision maker's decision set is a determining factor for the optimal learning rate. These results hold in a completely adversarial setting.

## 1 Introduction

Online convex optimization has emerged as a popular approach to online learning, bringing together convex optimization methods to tackle problems where repeated decisions need to be made in an unknown, possibly adversarial, environment. A full-information online convex optimization problem is a repeated zero-sum game between a learner (the player) and the environment (the opponent). There are  $T$  time periods. At each round  $t$ , the player has to choose  $f_t$  in a convex set  $\mathcal{F}$ . Subsequent to the choice of  $f_t$ , the environment reveals  $z_t \in \mathcal{Z}$  and the loss incurred to the player is  $\ell(z_t, f_t)$ , for a loss function  $\ell$  that is convex in its second argument. Both players are aware of all the parameters of the game, namely  $\ell$ ,  $\mathcal{Z}$ , and  $\mathcal{F}$ , prior to starting the game. Additionally, at the end of each period, the opponent's move is revealed to the player. The performance of the player is measured in terms of a quantity coined regret, defined as the gap between the accumulated losses incurred by the player and the best performance he could have achieved in hindsight with a non-adaptive strategy:

$$r_T((z_t)_{t=1,\dots,T}, (f_t)_{t=1,\dots,T}) = \sum_{t=1}^T \ell(z_t, f_t) - \inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(z_t, f).$$

In this field, one of the primary focus is to design algorithms, i.e., strategies to select  $(f_t)_{t=1,\dots,T}$  so as to keep the regret as small as possible even when facing an adversarial opponent. Particular

emphasis is placed on how the regret scales with  $T$  because this dependence relates to a notion of learning rate. If  $r_T = o(T)$ , the player is, in some sense, learning the game in the long-run since the gap between experienced and best achievable average cumulative payoffs vanishes as  $T \rightarrow \infty$ . Furthermore, the smaller the growth rate of  $r_T$ , the faster the learning. A natural question to ask is what is the best learning rate that can be achieved for a given game  $(\ell, \mathcal{Z}, \mathcal{F})$ . Mathematically, this is equivalent to characterizing the growth rate of the smallest regret that can be achieved by a player against a completely adversarial opponent, expressed as:

$$R_T(\ell, \mathcal{Z}, \mathcal{F}) = \inf_{f_1 \in \mathcal{F}} \sup_{z_1 \in \mathcal{Z}} \cdots \inf_{f_T \in \mathcal{F}} \sup_{z_T \in \mathcal{Z}} \left[ \sum_{t=1}^T \ell(z_t, f_t) - \inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(z_t, f) \right], \quad (1)$$

which we refer to as the value of the game  $(\ell, \mathcal{Z}, \mathcal{F})$ . Aside from pure learning considerations, the growth rate of  $R_T(\ell, \mathcal{Z}, \mathcal{F})$  has important consequences in a variety of fields where no-regret algorithms are used to compute complex quantities, e.g. Nash equilibria in Game Theory [14] or solutions to optimization problems in convex optimization [7], in which case this growth rate translates into the number of iterations required to compute the quantity with a given precision. We investigate this question in a systematic fashion by looking at the interplay between  $\mathcal{F}$  and  $\mathcal{Z}$  for the following class of piecewise linear loss functions:

$$\ell(z, f) = \max_{x \in \mathcal{X}(z)} (C(z)f + c(z))^T x, \quad (2)$$

where, for any  $z \in \mathcal{Z}$ ,  $C(z)$  is a matrix,  $c(z)$  is a vector, and  $\mathcal{X}(z) \subset \mathbb{R}^d$  is either a finite set or a polyhedron  $\{x \in \mathbb{R}^d \mid A(z)x \leq b(z)\}$  with  $A(z)$  a matrix and  $b(z)$  a vector. This type of loss functions arises in a number of important contexts such as online linear optimization, repeated zero-sum Stackelberg games, online prediction with side information, and online two-stage optimization, as illustrated in Section 1.1. Throughout the paper, we make the following standard assumption so as to have the game well-defined.

**Assumption 1**  $\mathcal{Z}$  is a non-empty compact subset of  $\mathbb{R}^{d_z}$  and  $\mathcal{F}$  is a non-empty, convex, and compact subset of  $\mathbb{R}^{d_f}$ . For any choice of  $z \in \mathcal{Z}$ , the set  $\mathcal{X}(z)$  is not empty. The loss function  $\ell$  is bounded on  $\mathcal{Z} \times \mathcal{F}$ . Moreover, either  $\mathcal{Z}$  has finite cardinality or  $\ell(\cdot, f)$  is continuous for any  $f \in \mathcal{F}$ .

**Contributions.** A number of no-regret algorithms developed in the literature can be used as a black box for the settings considered in this paper in order to get  $O(\sqrt{T})$  bounds on regret, e.g. Exponential Weights [20], Online Gradient Descent [22], and more generally Online Mirror Descent [9], to cite a few. To get better learning rates, other approaches have been proposed but they usually rely on either the curvature of  $\ell$ , for instance if  $\ell$  is strongly convex in its second argument [8], which is not the case here, or more information about the sequence  $(z_t)_{t=1, \dots, T}$ , see for example [17], which is not available in the fully adversarial setting. Aside from particular instances, e.g. [6] and [1], it is in general unknown how the interplay between  $\ell$ ,  $\mathcal{Z}$ , and  $\mathcal{F}$  determines the growth rate of  $R_T(\ell, \mathcal{Z}, \mathcal{F})$ . The main insight from this paper is that the curvature of the decision maker's set of moves is a determining factor for the growth rate of  $R_T(\ell, \mathcal{Z}, \mathcal{F})$ : if  $\mathcal{F}$  has rough edges then we are doomed to a rate of  $\Theta(\sqrt{T})$ , otherwise, if it is curved, the rate can be exponentially smaller. Specifically, we show that:

1. When  $\mathcal{F}$  is a polyhedron, either  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = 0$  or  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = \Omega(\sqrt{T})$ . This lower bound applies to online combinatorial optimization where  $\mathcal{F}$  is a combinatorial set, to many experts settings and repeated zero-sum Stackelberg games where the player resorts to a randomized strategy, as well as to many online prediction problems with side information and online two-stage optimization problems.
2. When (i)  $\ell$  is linear, (ii)  $\mathcal{F} = \{f \in \mathbb{R}^{d_f} \mid F(f) \leq 0\}$ , for  $F$  either a strongly convex function or  $F(f) = \|f\|_{\mathcal{F}} - C$  where  $C \geq 0$  and  $\|\cdot\|_{\mathcal{F}}$  is a  $q$ -uniformly convex norm with  $q \in [2, 3]$ , and (iii) 0 does not lie in the convex hull of  $\mathcal{Z}$ , we have  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = o(\sqrt{T})$ , achieved by the Follow-The-Leader algorithm [12]. This result applies to repeated zero-sum games where the player picks a cost vector (e.g. arc costs) of bounded euclidean norm and the opponent chooses an element in a combinatorial set (e.g. a path). This also applies to non-linear loss functions when 0 does not lie in the convex hull of the set of subgradients of  $\ell$  with respect to the second-coordinate by a standard reduction to linear loss functions,

see [22]. Note that assumption (iii) is required to get  $o(\sqrt{T})$  rates as it is well known that  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = \Omega(\sqrt{T})$  for linear losses when 0 lies in the interior of the convex hull of  $\mathcal{Z}$ , see Section 2.

## 1.1 Applications

We list examples of situations where losses of the type (2) arise.

**Online linear optimization** In this setting, the loss function is given by  $\ell(z, f) = z^\top f$ . This includes, in particular:

- online combinatorial optimization where the opponent picks a cost in  $[0, 1]^{d_z}$  and  $\mathcal{F}$  is defined as the convex hull of a finite set of elements (e.g. paths, spanning trees, and matchings),
- experts settings where the player picks a distribution over the experts' advice (in which case  $\mathcal{F}$  is also a polyhedron) and the opponent reveals a cost for each of the experts.

In online linear optimization, regret lower bounds are often derived by introducing a randomized zero-mean i.i.d. opponent, see [1]. However, this is possible only if 0 is in interior of the convex hull of  $\mathcal{Z}$ , which is typically not the case in online combinatorial optimization. A general feature of online linear optimization that will turn out to be important in the analysis is that there is no loss of generality in assuming that  $\mathcal{Z}$  is a convex set in the following sense.

**Lemma 1** *When  $\ell(z, f) = z^\top f$ , the games  $(\ell, \mathcal{Z}, \mathcal{F})$  and  $(\ell, \text{conv}(\mathcal{Z}), \mathcal{F})$  are equivalent, i.e.:*

$$R_T(\ell, \mathcal{Z}, \mathcal{F}) = R_T(\ell, \text{conv}(\mathcal{Z}), \mathcal{F}).$$

**Repeated zero-sum Stackelberg games** A repeated zero-sum Stackelberg game is a repeated zero-sum game with the particularity that one of the players, referred to as the leader, has to commit first to a randomized strategy  $f$  without even knowing which of the  $N$  other players, indexed by  $z$ , he is going to face at the next round. The interaction between the leader and player  $z \in \{1, \dots, N\}$  is captured by a payoff matrix  $M(z)$ . Once the leader is set on a strategy, the identity of the other player is revealed and the latter best-responds to the leader's strategy, leading to the following expression for the loss function:

$$\ell(z, f) = \max_{i=1, \dots, I_z} e_i^\top M(z) f,$$

where  $I_z$  is the number of possible moves for player  $z$ . We illustrate with a network security problem that has applications in urban network security [10] and fare evasion prevention in transit networks [11]. Consider a directed graph  $G = (V, E)$ . The leader has a limited number of patrols that can be assigned to arcs in order to intercept the attackers. A configuration  $\gamma \in \Gamma$  corresponds to a valid assignment of patrols to arcs and is represented by a vector  $(Y_{ij}^\gamma)_{(i,j) \in E}$  with  $Y_{ij}^\gamma = 1$  if a patrol is assigned to arc  $(i, j)$  and  $Y_{ij}^\gamma = 0$  otherwise. The leader chooses a mixed strategy  $f$  over the set of feasible allocations. Attacker  $z \in \{(i_1, j_1), \dots, (i_N, j_N)\}$  wants to go from  $z_1$  to  $z_2$  while minimizing the probability of being intercepted. This interaction is captured by the loss function:

$$\ell(z, f) = \max_{x \in \mathcal{X}(z)} \sum_{\gamma \in \Gamma} -f_\gamma x_\gamma,$$

with:

$$\mathcal{X}(z) = \left\{ \left( \max_{(i,j) \in E} X_{ij}^\pi Y_{ij}^\gamma \right)_{\gamma \in \Gamma} \mid \pi \in \Pi(z) \right\},$$

where  $\Pi(z)$  is the set of directed paths joining  $z_1$  to  $z_2$  in  $G$  and  $X_{ij}^\pi = 1$  if  $(i, j) \in \pi$  and  $X_{ij}^\pi = 0$  otherwise. The presentation of repeated Stackelberg games given here follows the model introduced by Balcan et al. [3] for general, i.e. not necessarily zero-sum, Stackelberg security games. In this more general setting, the loss function may not be convex and a possible approach, see [3], is to add another layer of randomization which casts the problem back into the realm of online linear optimization.

**Online prediction with side-information** This setting has a slightly different flavor as the opponent provides some side information  $x$  before the player gets to pick  $f \in \mathcal{F}$ , subsequent to what the opponent reveals the correct prediction  $y$ . Nonetheless, the lower bounds established in this paper also apply to this setting through a reduction to the setting without side-information, as detailed at the end of Section 2. In the standard linear binary prediction problem, where  $\mathcal{F}$  is a  $L^2$  ball,  $y \in \{-1, 1\}$ , and  $x$  lies in a  $L^2$  ball, loss functions of the form (2) are commonly used, e.g. the absolute loss  $\ell((x, y), f) = |y - x^\top f|$  and the hinge loss  $\ell((x, y), f) = \max(0, 1 - yx^\top f)$ . This is also true for linear multiclass prediction problems with the multiclass hinge loss:

$$\ell(f, (x, y)) = \max_{j=1, \dots, N} (1\{j \neq y\} + f_j^\top x - f_y^\top x),$$

where  $N$  denotes the number of classes,  $y \in \{1, \dots, N\}$ , and  $f$  is a vector obtained by concatenation of the vectors  $f_1, \dots, f_N$ . In the online approach to collaborative filtering, a typical loss function is  $\ell(M, (i, j, y)) = |M(i, j) - y|$  where  $M$  is a (user, item) matrix with bounded trace norm,  $(i, j)$  is a (user, item) pair, and  $y$  is the rating of item  $j$  by user  $i$ .

**Online two-stage optimization** This setting captures situations where the decision making process can be broken down into two consecutive stages. In the first stage, the player makes a decision represented by  $f \in \mathcal{F}$ . Subsequently, the opponent discloses some information  $z \in \mathcal{Z}$ , e.g. a demand vector, and then the player chooses another decision vector  $x$  in the second stage, taking into account this newly available information to optimize his objective function. The loss function takes on the following form:

$$\ell(z, f) = c_1^\top f + \min_{\substack{x \in \mathbb{R}^d \\ Af + Bx \leq z}} c_2^\top x,$$

where  $c_1$  and  $c_2$  are cost vectors and  $A$  and  $B$  are matrices. Using strong duality, this loss function can be expressed in the canonical form (2). This framework finds applications in the operation of power grids, where  $z$  represents the demand in electricity or the availability of various energy sources. Since  $z$  is unknown when it is time to set up conventional generators, the decision maker has to adjust the production or buy additional capacity from a spot market to meet the demand, see for example [13].

**Congestion control** We consider a variant of the congestion network game described in [4]. A decision maker has to decide how to ship a given set commodities through a network  $G = (V, E)$ . His decision can be equivalently represented by a flow vector  $f$ . Because the amount of commodities is assumed to be substantial, implementing  $f$  will cause congestion which will impact the other users of the network, represented by a flow vector  $z$ . The problem faced by the decision maker is to cause as little delay as possible to the other users with the additional difficulty that the traffic pattern  $z$  is not known ahead of time. Each arc  $e \in E$  has an associated latency function that is convex in the flow on this arc:

$$c_e(f + z) = \max_{k=1, \dots, K} (c_e^k \cdot (f_e + z_e) + s_e^k),$$

As a result, the total delay incurred to the other users can be expressed as:

$$\ell(z, f) = \sum_{e \in E} z_e \max_{k=1, \dots, K} (c_e^k \cdot (f_e + z_e) + s_e^k).$$

## 1.2 Related work

Asymptotically matching lower and upper bounds on  $R_T(\ell, \mathcal{Z}, \mathcal{F})$  can be found in the literature for a variety of loss functions although the discussion tends to be restrictive as far as the decision sets  $\mathcal{F}$  and  $\mathcal{Z}$  are concerned. The value of the game is shown to be  $\Theta(\log T)$  for three standard examples of curved loss functions. The first example, studied by Abernethy et al. [1], is the quadratic loss where  $\ell(z, f) = z \cdot f + \sigma \|f\|_2^2$  for  $\sigma > 0$ , with  $\mathcal{Z}$  and  $\mathcal{F}$  bounded  $L^2$  balls. The second, studied by Vovk [21], is the online linear regression setting where the opponent plays  $z = (y, x) \in \mathcal{Z} = [-C_y, C_y] \times B_\infty(0, 1)$  for  $C_y > 0$  ( $B_\infty(0, 1)$  denotes the unit  $L^\infty$  ball), the loss is  $\ell((y, x), f) = (y - x \cdot f)^2$ , and  $\mathcal{F}$  is an  $L^2$  ball. The last one, from Ordentlich and Cover [15], is the log-loss  $\ell(z, f) = -\log z \cdot f$  with  $\mathcal{Z}$  any compact set in  $\mathbb{R}^d$  and  $\mathcal{F}$  the simplex of dimension  $d$ . For non-curved losses, evidence suggests that the value of the game increases exponentially to  $\Theta(\sqrt{T})$ . Indeed,  $\Omega(\sqrt{T})$  lower bounds are proved for several instances involving the absolute loss  $\ell(z, f) = |z - f|$  in [5], typically with

$\mathcal{Z} = \{0, 1\}$  and  $\mathcal{F} = [0, 1]$ . For purely linear loss functions, Abernethy et al. [1] establish a  $\Omega(\sqrt{T})$  lower bound on  $R_T(\ell, \mathcal{Z}, \mathcal{F})$  when  $\mathcal{Z}$  is an  $L^2$  ball centered at 0 and  $\mathcal{F}$  is either an  $L^2$  ball or a bounded rectangle. This result was later generalized in [2] and shown to hold for  $\mathcal{F}$  a unit ball in any norm centered at 0 and  $\mathcal{Z}$  its dual ball. Cesa-Bianchi and Lugosi [6] investigate the experts setting, i.e.  $\mathcal{Z} = [0, 1]^d$  and  $\mathcal{F}$  is the simplex in dimension  $d$ , and proves the same  $\Omega(\sqrt{T})$  lower bound (which also holds if  $\mathcal{Z}$  is the simplex in dimension  $d$ , see [2]). Rakhlin et al. [19] establish  $\Omega(\sqrt{T})$  lower bounds on regret when  $\ell$  is the absolute loss for a prediction with side-information setting more general than the one considered in this paper where the player picks a function  $f(\cdot)$ , the opponent picks a pair  $(x, y)$ , and the loss is  $\ell(f(\cdot), (x, y)) = |f(x) - y|$ . The list of results listed above is far from being exhaustive but provides a good picture of the current state of the art. For each loss function, the intrinsic limitations of online algorithms are well-understood, usually with the construction of a particular example of  $\mathcal{F}$  and  $\mathcal{Z}$  for which a lower bound on  $R_T(\ell, \mathcal{Z}, \mathcal{F})$  asymptotically matches the best guarantee achieved by one of these algorithms. We aim at studying lower bounds on the value of the game in a more systematic fashion using tools rooted in duality theory and sensitivity analysis. All the proofs are deferred to the Appendix.

**Notations** For a set  $S \subset \mathbb{R}^d$ ,  $\text{conv}(S)$  (resp.  $\text{int}(S)$ ) refers to the convex hull (resp interior) of this set. When  $S$  is a compact, we define  $\mathcal{P}(S)$  as the set of probability measures on  $S$ . For  $x \in \mathbb{R}^d$ ,  $\|x\|$  refers to the  $L^2$  norm of  $x$  while  $B_2(x, \epsilon)$  denotes the closed  $L^2$  ball centered at  $x$  with width  $\epsilon$ . For a collection of random variables  $(Z_1, \dots, Z_t)$ ,  $\sigma(Z_1, \dots, Z_t)$  refers to the sigma-field generated by  $Z_1, \dots, Z_t$ . For a random variable  $Z$  and a probability distribution  $p$ , we write  $Z \sim p$  if  $Z$  is distributed according to  $p$ .

## 2 Lower bounds

Unless otherwise stated, we assume throughout this section that  $\ell$  can be written in the form (2). In particular,  $\ell(z, \cdot)$  is continuous for any  $z \in \mathcal{Z}$ . We build on a powerful result rooted in von Neumann's minimax theorem that enables the derivation of tight lower and upper bounds on  $R_T(\ell, \mathcal{Z}, \mathcal{F})$  by recasting the value of the game in a backward order.

**Theorem 1** *From [2]*

$$R_T(\ell, \mathcal{Z}, \mathcal{F}) = \sup_p \mathbb{E} \left[ \sum_{t=1}^T \inf_{f_t \in \mathcal{F}} \mathbb{E}[\ell(Z_t, f_t) | Z_1, \dots, Z_{t-1}] - \inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(Z_t, f) \right],$$

where the supremum is taken over the distribution  $p$  of the random variable  $(Z_1, \dots, Z_T)$  in  $\mathcal{Z}^T$ .

Any choice for  $p$  yields a lower bound on  $R_T(\ell, \mathcal{Z}, \mathcal{F})$ . The following result identifies a canonical choice for  $p$  that leads to  $\Omega(\sqrt{T})$  lower bounds on regret.

**Lemma 2** *Adapted from [2]*

*If we can find a distribution  $p$  on  $\mathcal{Z}$  and two points  $f_1$  and  $f_2$  in  $\arg\min_{f \in \mathcal{F}} \mathbb{E}[\ell(Z, f)]$  such that  $\ell(Z, f_1) \neq \ell(Z, f_2)$  with positive probability for  $Z \sim p$ , then  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = \Omega(\sqrt{T})$ .*

A distribution  $p$  satisfying the requirements of Lemma 2 can be viewed as an equalizing strategy for the opponent. This concept, formalized in [18], roughly refers to randomized strategies played by the opponent that cause the player's decisions to be completely irrelevant from a regret standpoint. These strategies are intrinsically hard to play against and often lead to tight lower bounds. To gain some intuition about this result, suppose that the opponent generates an independent copy of  $Z$  at each round  $t$ , which we denote by  $Z_t$ . In the adversarial setting considered in this paper, the player is aware of the opponent's strategy but does not get to see the realization of  $Z_t$  before committing to a decision. For this reason, at any round,  $f_1$  and  $f_2$  are optimal moves that are completely equivalent from the player's perspective. However, in hindsight, i.e. once all the realizations of the  $Z_t$ 's have been revealed,  $f_1$  and  $f_2$  are typically not equivalent because  $\ell(Z_t, f_1) \neq \ell(Z_t, f_2)$  with positive probability and one of these two moves will turn out to be

$$\max(0, \sum_{t=1}^T \ell(Z_t, f_1) - \ell(Z_t, f_2))$$

suboptimal which, in expectations, is of order  $\Omega(\sqrt{T})$  by the central limit theorem. Given the conditions imposed on  $p$ , it is convenient to work with the following equivalence relation.

**Definition 1** We define the equivalence relation  $\sim_\ell$  on  $\mathcal{F}$  by  $f_a \sim_\ell f_b$  for  $f_a, f_b \in \mathcal{F}$  if and only if  $\ell(z, f_a) = \ell(z, f_b)$  for all  $z \in \mathcal{Z}$ .

In what follows, we show, using sensitivity analysis for linear programming, that we can systematically, with the only exception of trivial games defined below, construct a distribution  $p$  with support  $\mathcal{Z}$  such that there are at least two equivalent classes in  $\operatorname{argmin}_{f \in \mathcal{F}} \mathbb{E}[\ell(Z, f)]$  whenever  $\mathcal{F}$  is a polyhedron, for  $Z \sim p$ .

**Definition 2** The game  $(\ell, \mathcal{Z}, \mathcal{F})$  is said to be trivial if and only if

$$\exists f^* \in \mathcal{F} \text{ such that } \forall z \in \mathcal{Z}, \ell(z, f^*) \leq \min_{f \in \mathcal{F}} \ell(z, f).$$

A simple example of a trivial game is  $(\ell(z, f) = zf, [0, 1], [0, 1])$ , where  $\ell(z, f) \geq 0 \forall f \in [0, 1]$  and  $\forall z \in [0, 1]$ , with  $\ell(z, f) = 0$  if  $f = 0$  irrespective of  $z$ . If the game is trivial, the player will always play  $f^*$  irrespective of the time horizon and of the opponent's strategy observed so far to obtain zero regret. As it turns out, this uniquely identifies trivial games as we establish in Lemma 3.

**Lemma 3** For any  $T \in \mathbb{N}$ ,  $R_T(\ell, \mathcal{Z}, \mathcal{F}) \geq 0$ . Moreover in any of the following cases:

1.  $\mathcal{Z}$  has finite cardinality,
2.  $\ell(\cdot, f)$  is continuous for any choice of  $f \in \mathcal{F}$ ,

$R_T(\ell, \mathcal{Z}, \mathcal{F}) = 0$  if and only if the game is trivial.

The following result shows that, in most cases of interest, we can drastically restrict the power of the opponent while still preserving the nature of the game. This enables us to focus on the case where  $\mathcal{Z}$  is finite.

**Lemma 4** Suppose that  $\ell(\cdot, f)$  is continuous for any choice of  $f \in \mathcal{F}$ . If the game  $(\ell, \mathcal{Z}, \mathcal{F})$  is not trivial, there exists a finite subset  $\tilde{\mathcal{Z}} \subseteq \mathcal{Z}$  such that the game  $(\ell, \tilde{\mathcal{Z}}, \mathcal{F})$  is not trivial.

We are now ready to present the main results of this section. To the best of our knowledge, these results constitute the first systematic  $\Omega(\sqrt{T})$  lower bounds on regret obtained for a large class of piecewise linear loss functions.

**Theorem 2** Suppose that  $\mathcal{F}$  is a polyhedron. In any of the following cases:

1.  $\mathcal{Z}$  has finite cardinality,
2.  $\ell(\cdot, f)$  is continuous for any choice of  $f \in \mathcal{F}$ ,

either the game is trivial or  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = \Omega(\sqrt{T})$ .

An immediate consequence of Theorem 2 for linear games is the following:

**Theorem 3** Suppose that  $\mathcal{F}$  is a polyhedron and that  $\ell(z, f) = z^\top f$ . Then, either the game is trivial or  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = \Omega(\sqrt{T})$ .

The proofs rely on Lemma 2 which is based on Theorem 1 and may, as a result, seem rather obscure. We stress that these lower bounds are derived by means of an equalizing strategy. We present this more intuitive view in the Appendix by exhibiting an equalizing strategy in the online linear optimization setting when  $\operatorname{int}(\operatorname{conv}(\mathcal{Z})) \neq \emptyset$ .

Note that Theorems 2 and 3 imply  $\Omega(\sqrt{T})$  regret for a number of repeated Stackelberg games and online linear optimization problems as discussed in Section 1.1. Furthermore, we stress that Theorem 2 can also be used when  $\mathcal{F}$  is not a polyhedron but this typically requires a preliminary step which boils down to restricting the opponent's decision set. For instance, the following well-known result is almost a direct consequence of Theorem 3.

**Lemma 5** Suppose that  $\ell(z, f) = z^\top f$ , that  $0 \in \text{int}(\text{conv}(\mathcal{Z}))$ , and that  $\mathcal{F}$  contains at least two elements. Then  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = \Omega(\sqrt{T})$ .

Note that Lemma 5 is consistent with Theorem 3 as the game  $(\ell(z, f) = z^\top f, \mathcal{Z}, \mathcal{F})$  is non-trivial if  $0 \in \text{int}(\text{conv}(\mathcal{Z}))$  as soon as  $\mathcal{F}$  contains at least two elements. Indeed:

$$\ell(\epsilon \frac{f_2 - f_1}{\|f_2 - f_1\|}, f_2) > \ell(\epsilon \frac{f_2 - f_1}{\|f_2 - f_1\|}, f_1) \text{ and } \ell(\epsilon \frac{f_1 - f_2}{\|f_2 - f_1\|}, f_1) > \ell(\epsilon \frac{f_1 - f_2}{\|f_2 - f_1\|}, f_2),$$

for a small enough  $\epsilon > 0$  and any pair  $f_1 \neq f_2 \in \mathcal{F}$ . When  $0 \in \text{int}(\mathcal{Z})$ , the opponent has some freedom to play, at each time period, a random vector with expected value zero, making every strategy available to the player equally bad. In other words, any i.i.d. zero-mean distribution is an equalizing strategy for the opponent in this case.

A preliminary step is also required to derive  $\Omega(\sqrt{T})$  lower bounds on regret for prediction problems with side information where  $\mathcal{F}$  is typically not a polyhedron. We sketch this simple argument for the canonical classification problem with the hinge loss, i.e. the game :

$$(\ell((x, y), f) = \max(0, 1 - yx^\top f), \mathcal{Z} = B_2(0, 1) \times \{-1, 1\}, \mathcal{F} = B_2(0, 1)),$$

but the method readily extends to any of the prediction problems described in Section 1.1. The idea is to restrict the opponent's decision set by taking a fixed vector  $x$  of norm 1 and impose that, at any round  $t$ , the opponent's move be  $(x, y_t)$  for  $y_t \in \{-1, 1\}$ . Since  $\ell((x, y), f)$  only depends on  $f$  through the scalar product between  $f$  and  $x$ , the player's decision set can be equivalently described by this value, which lies in  $[-1, 1]$ . Formally, we define a new loss function  $\tilde{\ell}(y, f) = \max(0, 1 - yf)$  with  $\tilde{\mathcal{Z}} = \{-1, 1\}$  and  $\tilde{\mathcal{F}} = [-1, 1]$  and we have:

$$R_T(\ell, \mathcal{Z}, \mathcal{F}) \geq R_T(\tilde{\ell}, \tilde{\mathcal{Z}}, \tilde{\mathcal{F}}).$$

Observe now that the game  $(\tilde{\ell}, \tilde{\mathcal{Z}}, \tilde{\mathcal{F}})$  is not trivial, that  $\tilde{\mathcal{Z}}$  is discrete, and that:

$$\tilde{\ell}(y, f) = \max_{\alpha \in \{0, 1\}} (\alpha, -y\alpha)^\top (1, f).$$

We conclude with Theorem 2 that  $R_T(\tilde{\ell}, \tilde{\mathcal{Z}}, \tilde{\mathcal{F}}) = \Omega(\sqrt{T})$  which implies that  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = \Omega(\sqrt{T})$ .

**Remark about Lemma 2** We point out that, in general, it is not possible to weaken the assumptions of Lemma 2 (which, in fact, applies to a much more general class of loss functions than the one given by (2)). In particular, finding  $z \in \mathcal{Z}$  such that there are two equivalence classes  $f_1$  and  $f_2$  in  $\text{argmin}_{f \in \mathcal{F}} \ell(z, f)$  does not guarantee that  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = \Omega(\sqrt{T})$  as we illustrate with a counterexample. This is because the result of Lemma 2 is intrinsically tied to the central limit theorem. Consider the following (non-trivial) online linear regression game:

$$(\ell(z, f) = (z^\top f)^2, \mathcal{Z} = B_2(z^*, 1), \mathcal{F} = [f_1, f_2]),$$

where  $f_1 = (1, 0, \dots, 0)$ ,  $f_2 = (0, 1, 0, \dots)$ , and  $z^* = (1, 1, 0, \dots, 0)$ . Observe that  $\forall z \in \mathcal{Z}, \forall f \in \mathcal{F}, z^\top f \geq 0$ . Hence  $\text{argmin}_{f \in \mathcal{F}} \ell(z, f) = \text{argmin}_{f \in \mathcal{F}} f^\top z$ . Furthermore,  $\text{argmin}_{f \in \mathcal{F}} f^\top z^* = [f_1, f_2]$  but  $f_1$  and  $f_2$  are clearly not in the same equivalence class. Yet, even though  $\ell$  is not strongly convex in  $f$ , there exists an algorithm achieving  $O(\log(T))$  regret, see [21].

### 3 Upper bounds

Looking at Theorem 2, Theorem 3, and Lemma 5, it is tempting to conclude that the existence of pieces where  $\ell$  is linear in its second argument dooms us to  $\Omega(\sqrt{T})$  regret bounds. We argue that the growth rate of  $R_T(\ell, \mathcal{Z}, \mathcal{F})$  is determined by a more involved interplay between  $\ell$ ,  $\mathcal{Z}$ , and  $\mathcal{F}$  so that this assertion requires further examination. In fact, we show that  $O(\log(T))$  regret bounds are even possible in online linear optimization. The fundamental reason is that the curvature of the boundary of  $\mathcal{F}$  can make up for the lack of thereof of  $\ell$ . Curvature is key to enforce stability of the player's strategy with respect to perturbations in the opponent's moves. Sometimes, when the predictions are stable, e.g. when  $\ell$  is the square loss  $\ell(z, f) = \|z - f\|^2$ , a very simple algorithm, known as Follow-The-Leader, yields  $O(\log(T))$  regret.

**Definition 3** From [12]

The Follow-The-Leader (FTL) strategy consists in playing:

$$f_t \in \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{t-1} \sum_{\tau=1}^{t-1} \ell(z_\tau, f).$$

It is well-known that FTL fails to yield sublinear regret for online linear optimization in general. However, when  $\mathcal{F}$  has a strongly curved boundary and  $0 \notin \operatorname{conv}(\mathcal{Z})$ , the FTL strategy becomes stable, leading to  $O(\log(T))$  regret as we next show using sensitivity analysis.

**Theorem 4** Suppose that (i)  $\ell$  is linear, (ii)  $\mathcal{F} = \{f \in \mathbb{R}^{d_f} \mid F(f) \leq 0\}$  for  $F$  a strongly convex function with respect to the  $L^2$  norm, and (iii)  $0 \notin \operatorname{conv}(\mathcal{Z})$ . Then, FTL yields  $O(\log(T))$  regret.

As an example of application of Theorem 4, consider a repeated network game where the player picks the arc costs in a  $L^2$  ball, the opponent picks a path, and the loss incurred to the opponent is the sum of the arc costs along the path. In this setting, FTL yields  $O(\log(T))$  regret even though the game is not trivial. Theorem 4 also has some implications for non-linear convex loss functions when the boundary of  $\mathcal{F}$  is curved and 0 is not in the convex hull of the set of subgradients of  $\ell$  with respect to the player's moves. Indeed suppose that, at any time period  $t$ , the player follows the FTL strategy as if the loss function were linear and the past moves of the opponents were  $y_1, \dots, y_{t-1}$ , i.e.:

$$f_t \in \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{t-1} \sum_{\tau=1}^{t-1} y_\tau^\top f,$$

where, for any  $\tau = 1, \dots, t-1$ ,  $y_\tau$  is a subgradient of  $\ell(z_t, \cdot)$  at  $f_t$ . Then, for any sequence of moves  $(z_1, \dots, z_T)$ , we have:

$$\begin{aligned} r_T((z_t)_{t=1, \dots, T}, (f_t)_{t=1, \dots, T}) &\leq \sum_{t=1}^T y_t^\top f_t - \inf_{f \in \mathcal{F}} \sum_{t=1}^T y_t^\top f \\ &= O(\log(T)). \end{aligned}$$

It is, however, a priori unclear whether  $\log(T)$  is the optimal growth rate for games satisfying the assumptions of Theorem 4. Quite surprisingly, i.i.d. opponents appear to be particularly weak for this kind of games, incurring at most a  $O(1)$  regret lower bound as shown in the following lemma. This is in stark contrast with the situations of Section 2 where the (tight)  $\Omega(\sqrt{T})$  lower bounds are always derived through i.i.d. opponents.

**Lemma 6** Consider the game  $(\ell(z, f) = z^\top f, \mathcal{Z}, \mathcal{F})$  with  $0 \notin \operatorname{conv}(\mathcal{Z})$  and  $\mathcal{F} = B_2(0, 1)$ . Any lower bound derived from Theorem 1 with i.i.d. random variables  $Z_1, \dots, Z_T \sim p$  is  $O(1)$  for any choice of  $p \in \mathcal{P}(\mathcal{Z})$ .

Abernethy et al. [2] remark that restricting the study to i.i.d. sequences is in general not enough to get tight bounds for non-linear losses such as  $\ell(z, f) = \|z - f\|^2$ . It turns out that this is also true for linear losses when  $\mathcal{F}$  is strongly curved and  $0 \notin \operatorname{conv}(\mathcal{Z})$  as the value of the game is in fact  $\Theta(\log(T))$ .

**Theorem 5** When  $\ell(z, f) = z^\top f$ ,  $0 \notin \operatorname{conv}(\mathcal{Z})$ , and  $\mathcal{F} = B_2(0, 1)$ , the value of the game is:

$$R_T(\ell, \mathcal{Z}, \mathcal{F}) = \Theta(\log(T)).$$

So far, we have studied two scenarios that are diametrically opposed in terms of the curvature of the decision sets with polyhedra on one side in Section 2, with  $\Theta(\sqrt{T})$  regret, and euclidean balls, with  $\Theta(\log(T))$  regret, on the other side of the spectrum. Interestingly, bridging this gap leads to the rise of intermediate learning rates that can be quantified through the modulus of convexity of  $\mathcal{F}$ . Specifically, consider any norm  $\|\cdot\|_{\mathcal{F}}$ . The modulus of convexity of the associated unit ball is defined as:

$$\delta_{\mathcal{F}} : \epsilon \rightarrow \inf_{\substack{\|f\|_{\mathcal{F}}, \|\tilde{f}\|_{\mathcal{F}} \leq 1 \\ \|f - \tilde{f}\|_{\mathcal{F}} \geq \epsilon}} 1 - \left\| \frac{f + \tilde{f}}{2} \right\|_{\mathcal{F}}.$$



The norm  $\|\cdot\|_{\mathcal{F}}$  is said to be uniformly convex if its characteristic of convexity is equal to 0, i.e.:

$$\sup\{\epsilon \geq 0 \mid \delta_{\mathcal{F}}(\epsilon) = 0\} = 0.$$

Pisier [16] show that if  $\|\cdot\|_{\mathcal{F}}$  is uniformly convex, there must exist  $q \geq 2$  and  $c > 0$  such that  $\delta_{\mathcal{F}}(\epsilon) \geq c\epsilon^q$  for all  $\epsilon \in [0, 2]$ , in which case we say that  $\|\cdot\|_{\mathcal{F}}$  is  $q$ -uniformly convex. This parameter quantifies how curved  $\|\cdot\|_{\mathcal{F}}$  balls are and determines the growth rate of the value of the game when  $\mathcal{F}$  is a  $\|\cdot\|_{\mathcal{F}}$  ball and  $0 \notin \text{conv}(\mathcal{Z})$ .

**Lemma 7** *Consider the game  $(\ell(z, f) = z^\top f, \mathcal{Z}, \mathcal{F})$  with  $0 \notin \text{conv}(\mathcal{Z})$  and  $\mathcal{F} = \{f \mid \|f\|_{\mathcal{F}} \leq C\}$ , where  $\|\cdot\|_{\mathcal{F}}$  is a  $q$ -uniformly convex norm and  $C \geq 0$ . Then, FTL yields regret  $O(\log(T))$  if  $q = 2$  and regret  $O(T^{\frac{q-2}{q-1}})$  if  $q \in (2, 3]$ .*

**Situations where 0 lies on the boundary of  $\mathcal{Z}$  when  $\ell$  is linear** Observe that this situation is not covered by Lemma 5, Theorem 4, Theorem 5, nor Lemma 7. We stress that zero-mean i.i.d. opponents are not helpful to derive  $\Omega(\sqrt{T})$  regret when 0 lies exactly on the boundary of  $\mathcal{Z}$  (in fact, in the relative interior of an edge of  $\mathcal{Z}$ ), as we illustrate with an example. Therefore, the growth rate of  $R_T(\ell, \mathcal{Z}, \mathcal{F})$  remains unknown in this setup. Define  $\mathcal{Z} = \text{conv}(z_1, z_2, z_3, z_4)$  with  $z_1 = (-1, 1, 0, 0)$ ,  $z_2 = (1, -1, 0, 0)$ ,  $z_3 = (0, 0, 0, 1)$ , and  $z_4 = (0, 0, 1, 0)$ . Also define  $\mathcal{F} = [f^*, f^{**}]$  with  $f^* = (0, 0, 0, 0)$  and  $f^{**} = (1, 1, -1, 1)$ . Observe that the game  $(\ell(z, f) = z^\top f, \mathcal{Z}, \mathcal{F})$  is not trivial because  $\arg\min_{f \in \mathcal{F}} f \cdot z_3 = \{f^*\}$  while  $\arg\min_{f \in \mathcal{F}} f \cdot z_4 = \{f^{**}\}$ . For any zero-mean i.i.d. opponent  $Z_1, \dots, Z_T$ , the only possibility is to have  $Z_t \in [z_1, z_2]$ . We get, irrespective of the player's strategy:

$$\mathbb{E}[r_T((Z_t)_{t=1, \dots, T}, (f_t)_{t=1, \dots, T})] = -\mathbb{E}[\inf_{f \in \mathcal{F}} f \cdot \sum_{t=1}^T Z_t].$$

We have  $f^* \in \arg\min_{f \in \mathcal{F}} f \cdot z$  for any  $z \in [z_1, z_2]$ . Hence:

$$\mathbb{E}[r_T((Z_t)_{t=1, \dots, T}, (f_t)_{t=1, \dots, T})] = -\mathbb{E}[f^* \cdot \sum_{t=1}^T Z_t],$$

which finally yields  $\mathbb{E}[r_T((Z_t)_{t=1, \dots, T}, (f_t)_{t=1, \dots, T})] = 0$ .

## Acknowledgments

The authors would like to thank Alexander Rakhlin for his valuable input, and in particular, for bringing to our attention the possibility of having  $o(\sqrt{T})$  bounds on regret in the linear setting.

## References

- [1] J. Abernethy, P. L. Bartlett, A. Rakhlin, and A. Tewari. Optimal strategies and minimax lower bounds for online convex games. In *Proc. of the 21st Annual Conf. on Learning Theory (COLT)*, pages 415–424, 2008.
- [2] J. Abernethy, A. Agarwal, P. L. Bartlett, and A. Rakhlin. A stochastic view of optimal regret through minimax duality. In *Proc. of the 22nd Annual Conf. on Learning Theory (COLT)*, 2009.
- [3] M. Balcan, A. Blum, N. Haghtalab, and A. D. Procaccia. Commitment without regrets: Online learning in stackelberg security games. In *Proc. 16th ACM Conf. on Economics and Computation (EC)*, pages 61–78, 2015.
- [4] V. Bonifaci, T. Harks, and G. Schäfer. Stackelberg routing in arbitrary networks. *Mathematics of Operations Research*, 35(2):330–346, 2010.
- [5] N. Cesa-Bianchi and G. Lugosi. On prediction of individual sequences. *Annals of Statistics*, 27(6):1865–1895, 1999.

- [6] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [7] Y. Freund and R. E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1):79–103, 1999.
- [8] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- [9] M. Herbster and M. K. Warmuth. Tracking the best expert. *Machine Learning*, 32(2):151–178, 1998.
- [10] M. Jain, V. Conitzer, and M. Tambe. Security scheduling for real-world networks. In *Proc. of the 2013 int. conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*, pages 215–222, 2013.
- [11] A. X. Jiang, Z. Yin, M. P. Johnson, M. Tambe, C. Kiekintveld, K. Leyton-Brown, and T. Sandholm. Towards optimal patrol strategies for fare inspection in transit systems. In *AAAI Spring Symposium: Game Theory for Security, Sustainability, and Health*, 2012.
- [12] A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- [13] S. Kim, G. B. Giannakis, and K. Y. Lee. Online optimal power flow with renewables. In *48th Asilomar Conf. on Signals, Systems and Computers*, pages 355–360. IEEE, 2014.
- [14] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic game theory*, volume 1. Cambridge University Press, 2007.
- [15] E. Ordentlich and T. M. Cover. The cost of achieving the best portfolio in hindsight. *Mathematics of Operations Research*, 23(4):960–982, 1998.
- [16] Gilles Pisier. Martingales with values in uniformly convex spaces. *Israel Journal of Mathematics*, 20(3-4):326–350, 1975.
- [17] A. Rakhlin and K. Sridharan. Optimization, learning, and games with predictable sequences. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3066–3074, 2013.
- [18] A. Rakhlin, K. Sridharan, and A. Tewari. Online learning: Beyond regret. In *Proc. of the 24th Annual Conf. on Learning Theory (COLT)*, pages 559–594, 2011.
- [19] A. Rakhlin, K. Sridharan, and A. Tewari. Online learning via sequential complexities. *The Journal of Machine Learning Research*, 16(1):155–186, 2015.
- [20] V. Vovk. Aggregating strategies. In *Proc. 3rd Workshop on Computational Learning Theory (COLT)*, pages 371–383, 1990.
- [21] V. Vovk. Competitive on-line linear regression. In *Advances in Neural Information Processing Systems (NIPS)*, pages 364–370, 1998.
- [22] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proc. of the 20th Int. Conf. on Machine Learning (ICML)*, pages 928–936, 2003.

## 4 Appendix: proofs

### 4.1 Proof of Theorem 1

The assumptions of Theorem 1 in [2] are satisfied for the game  $(\ell, \mathcal{Z}, \mathcal{F})$  using Assumption 1 and the fact that any loss function  $\ell$  of the form (2) is such that  $\ell(z, \cdot)$  is continuous for any  $z \in \mathcal{Z}$ .

## 4.2 Proof of Lemma 1

The proof follows from the repeated use of the Von Neumann's minimax theorem developed in [2]. To simplify the presentation, we prove the result when  $T = 2$  but the general proof follows the same principle. We have:

$$R_2 = \inf_{f_1 \in \mathcal{F}} \sup_{z_1 \in \mathcal{Z}} \inf_{f_2 \in \mathcal{F}} \sup_{z_2 \in \mathcal{Z}} \left[ \sum_{t=1}^2 \ell(z_t, f_t) - \inf_{f \in \mathcal{F}} \sum_{t=1}^2 \ell(z_t, f) \right].$$

Consider fixed vectors  $f_1, f_2 \in \mathcal{F}$  and  $z_1 \in \mathcal{Z}$  and define the function  $M(z_2) = \sum_{t=1}^2 \ell(z_t, f_t) - \inf_{f \in \mathcal{F}} \sum_{t=1}^2 \ell(z_t, f)$ . Observe that  $M$  is convex. Indeed,  $z_2 \rightarrow \ell(z_1, f_1) + \ell(z_2, f_2)$  is affine and  $z_2 \rightarrow \inf_{f \in \mathcal{F}} \sum_{t=1}^2 \ell(z_t, f)$  is concave as the infimum of affine functions. Therefore:

$$\sup_{z_2 \in \mathcal{Z}} M(z_2) = \sup_{z_2 \in \text{conv}(\mathcal{Z})} M(z_2).$$

We obtain:

$$R_2 = \inf_{f_1 \in \mathcal{F}} \sup_{z_1 \in \mathcal{Z}} \inf_{f_2 \in \mathcal{F}} \sup_{z_2 \in \text{conv}(\mathcal{Z})} \left[ \sum_{t=1}^2 \ell(z_t, f_t) - \inf_{f \in \mathcal{F}} \sum_{t=1}^2 \ell(z_t, f) \right].$$

By randomizing the choice of  $z_2$ , we can use Von Neumann's minimax theorem to derive:

$$R_2 = \inf_{f_1 \in \mathcal{F}} \sup_{z_1 \in \mathcal{Z}} \left\{ \ell(z_1, f_1) + \sup_{p_2 \in \mathcal{P}(\text{conv}(\mathcal{Z}))} \left\{ \inf_{f_2 \in \mathcal{F}} \mathbb{E}_{z \sim p_2} \ell(z, f_2) - \mathbb{E}_{z_2 \sim p_2} \inf_{f \in \mathcal{F}} \sum_{t=1}^2 \ell(z_t, f) \right\} \right\}.$$

For a fixed  $f_1 \in \mathcal{F}$ , define:

$$A(z_1) = \ell(z_1, f_1) + \sup_{p_2 \in \mathcal{P}(\text{conv}(\mathcal{Z}))} \left\{ \inf_{f_2 \in \mathcal{F}} \mathbb{E}_{z \sim p_2} \ell(z, f_2) - \mathbb{E}_{z_2 \sim p_2} \inf_{f \in \mathcal{F}} \sum_{t=1}^2 \ell(z_t, f) \right\}.$$

Observe that, for a fixed  $p_2 \in \mathcal{P}(\text{conv}(\mathcal{Z}))$ , the function:

$$z_1 \rightarrow \inf_{f_2 \in \mathcal{F}} \mathbb{E}_{z \sim p_2} \ell(z, f_2) - \mathbb{E}_{z_2 \sim p_2} \inf_{f \in \mathcal{F}} \sum_{t=1}^2 \ell(z_t, f)$$

is convex as the difference between a constant and the expected value of the infimum of affine functions. Since the supremum of convex functions is convex,  $A$  is convex and  $\sup_{z_1 \in \mathcal{Z}} A(z_1) = \sup_{z_1 \in \text{conv}(\mathcal{Z})} A(z_1)$ . We derive:

$$R_2 = \inf_{f_1 \in \mathcal{F}} \sup_{z_1 \in \text{conv}(\mathcal{Z})} \left[ \ell(z_1, f_1) + \sup_{p_2 \in \mathcal{P}(\text{conv}(\mathcal{Z}))} \left\{ \inf_{f_2 \in \mathcal{F}} \mathbb{E}_{z \sim p_2} \ell(z, f_2) - \mathbb{E}_{z_2 \sim p_2} \inf_{f \in \mathcal{F}} \sum_{t=1}^2 \ell(z_t, f) \right\} \right].$$

To conclude, we unwind the first step, i.e. we use the minimax theorem in reverse order. This yields:

$$R_2 = \inf_{f_1 \in \mathcal{F}} \sup_{z_1 \in \text{conv}(\mathcal{Z})} \inf_{f_2 \in \mathcal{F}} \sup_{z_2 \in \text{conv}(\mathcal{Z})} \left[ \sum_{t=1}^2 \ell(z_t, f_t) - \inf_{f \in \mathcal{F}} \sum_{t=1}^2 \ell(z_t, f) \right],$$

i.e.  $R_2(\ell(z, f) = z^\top f, \mathcal{Z}, \mathcal{F}) = R_2(\ell(z, f) = z^\top f, \text{conv}(\mathcal{Z}), \mathcal{F})$ . Moreover,  $\mathcal{Z}$  is a compact set which implies that  $\text{conv}(\mathcal{Z})$  is also a compact set by a standard topological argument. As a result, the game  $(\ell(z, f) = z^\top f, \text{conv}(\mathcal{Z}), \mathcal{F})$  also satisfies Assumption 1.

### 4.3 Proof of Lemma 2

We follow the analysis of Theorem 19 of [2]. Using Theorem 1 with  $p$  taken as the distribution of i.i.d. copies of  $Z$ , we get the lower bound:

$$\begin{aligned}
R_T &\geq T \inf_{f \in \mathcal{F}} \mathbb{E}[\ell(Z_t, f)] - \mathbb{E}[\inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(Z_t, f)] \\
&\geq T \sup_{f \in \{f_1, f_2\}} \mathbb{E}[\ell(Z_t, f)] - \mathbb{E}[\inf_{f \in \{f_1, f_2\}} \sum_{t=1}^T \ell(Z_t, f)] \\
&\geq \mathbb{E}[\max\{\sum_{t=1}^T \mathbb{E}[\ell(Z_t, f_1)] - \ell(Z_t, f_1), \sum_{t=1}^T \mathbb{E}[\ell(Z_t, f_2)] - \ell(Z_t, f_2)\}] \\
&\geq \mathbb{E}[\max\{0, \sum_{t=1}^T \ell(Z_t, f_1) - \ell(Z_t, f_2)\}],
\end{aligned}$$

where we use the fact that  $\inf_{f \in \mathcal{F}} \mathbb{E}[\ell(Z_t, f)] = \mathbb{E}[\ell(Z_t, f_1)] = \mathbb{E}[\ell(Z_t, f_2)]$ . Since  $\ell(Z, f_2) \neq \ell(Z, f_1)$  with positive probability, the random variables  $(\ell(Z_t, f_1) - \ell(Z_t, f_2))_{t=1, \dots, T}$  are i.i.d. with zero mean and positive variance and we can conclude with the central limit theorem since  $\ell$  is bounded.

### 4.4 Proof of Lemma 3

The fact that  $R_T \geq 0$  is proved in Lemma 3 of [2] and follows from Theorem 1 by taking the  $Z_t$ 's to be deterministic and all equal to any  $z \in \mathcal{Z}$ . Clearly, if the game is trivial then  $R_T = 0$  because this value is attained for  $f_1, \dots, f_T = f^*$  irrespective of the decisions made by the opponent. Conversely, suppose  $R_T = 0$ . Consider  $p$  to be the product of  $T$  uniform distributions on  $\mathcal{Z}$ . Then, using again Theorem 1:

$$0 \geq \mathbb{E}[\sum_{t=1}^T \inf_{f_t \in \mathcal{F}} \mathbb{E}[\ell(Z_t, f_t)] - \inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(Z_t, f)],$$

as  $Z_1, \dots, Z_T$  are independent random variables. Since they are also identically distributed, we obtain:

$$0 \geq T \cdot \inf_{f \in \mathcal{F}} \mathbb{E}[\ell(Z, f)] - \mathbb{E}[\inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(Z_t, f)].$$

Yet  $\mathbb{E}[\inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(Z_t, f)] \leq \inf_{f \in \mathcal{F}} \mathbb{E}[\sum_{t=1}^T \ell(Z_t, f)] = T \cdot \inf_{f \in \mathcal{F}} \mathbb{E}[\ell(Z, f)]$  and we derive:

$$T \cdot \inf_{f \in \mathcal{F}} \mathbb{E}[\ell(Z, f)] - \mathbb{E}[\inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(Z_t, f)] = 0.$$

Since  $\ell$  is bounded,  $\mathcal{Z}$  is compact, and  $\ell(z, \cdot)$  is continuous for any  $z \in \mathcal{Z}$ ,  $f \rightarrow \mathbb{E}[\ell(Z, f)]$  is continuous by dominated convergence so we can take  $f^* \in \operatorname{argmin}_{f \in \mathcal{F}} \mathbb{E}[\ell(Z, f)]$  ( $\mathcal{F}$  is compact). We obtain:

$$\mathbb{E}[\sum_{t=1}^T \ell(Z_t, f^*) - \inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(Z_t, f)] = 0.$$

As  $\sum_{t=1}^T \ell(Z_t, f^*) - \inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(Z_t, f) \geq 0$ , we derive that:

$$(z_1, \dots, z_T) \rightarrow \sum_{t=1}^T \ell(z_t, f^*) - \inf_{f \in \mathcal{F}} \sum_{t=1}^T \ell(z_t, f) = 0$$

almost everywhere on  $\mathcal{Z}^T$ . If  $\mathcal{Z}$  is discrete, this implies equality on  $\mathcal{Z}^T$ , which in particular implies  $\ell(z, f^*) = \inf_{f \in \mathcal{F}} \ell(z, f)$  for all  $z \in \mathcal{Z}$  and we are done. If, on the other hand,  $\ell(\cdot, f)$  is continuous for

all  $f \in \mathcal{F}$ , we have:

$$\sum_{t=1}^T \ell(z_t, f^*) \leq \sum_{t=1}^T \ell(z_t, f), \quad \forall f \in \mathcal{F}, \forall (z_1, \dots, z_T) \in \tilde{\mathcal{Z}},$$

for  $\tilde{\mathcal{Z}}$  a subset of  $\mathcal{Z}$  with Lebesgue measure equal to that of  $\mathcal{Z}$ . Since a non-empty open set cannot have Lebesgue measure 0,  $\tilde{\mathcal{Z}}$  is dense in  $\mathcal{Z}$  and by taking limits in the above inequality for each  $f \in \mathcal{F}$  separately, we conclude that:

$$\sum_{t=1}^T \ell(z_t, f^*) \leq \sum_{t=1}^T \ell(z_t, f), \quad \forall f \in \mathcal{F}, \forall (z_1, \dots, z_T) \in \mathcal{Z},$$

which in particular implies that  $\ell(z, f^*) = \inf_{f \in \mathcal{F}} \ell(z, f)$  for all  $z \in \mathcal{Z}$  and the game is trivial.

#### 4.5 Proof of Lemma 4

Suppose by contradiction that we cannot find such a finite subset. Since  $\mathcal{Z}$  is compact, it is also separable thus it contains a countable dense subset  $\{z_n \mid n \in \mathbb{N}\}$ . By assumption, the game  $(\ell, \{z_k \mid k \leq n\}, \mathcal{F})$  must be trivial for any  $n$ , i.e. there exists  $f_n \in \mathcal{F}$  such that:

$$\ell(z_k, f_n) \leq \min_{f \in \mathcal{F}} \ell(z, f), \quad \forall k \leq n.$$

Since  $\mathcal{F}$  is compact, we can find a subsequence of  $(f_n)_{n \in \mathbb{N}}$  such that  $f_n \rightarrow f^* \in \mathcal{F}$ . Without loss of generality, we continue to refer to this sequence as  $(f_n)_{n \in \mathbb{N}}$ . Taking the limit  $n \rightarrow \infty$  in the above inequality for any fixed  $k \in \mathbb{N}$  yields:

$$\ell(z_k, f^*) \leq \ell(z_k, f), \quad \forall f \in \mathcal{F}, \forall k \in \mathbb{N}.$$

Consider a fixed  $f \in \mathcal{F}$ , since  $\{z_n \mid n \in \mathbb{N}\}$  is dense in  $\mathcal{Z}$  and since  $\ell(\cdot, f^*)$  and  $\ell(\cdot, f)$  are continuous, we get:

$$\ell(z, f^*) \leq \ell(z, f), \quad \forall f \in \mathcal{F}, \forall z \in \mathcal{Z},$$

which shows that  $(\ell, \mathcal{Z}, \mathcal{F})$  is trivial, a contradiction.

#### 4.6 Proof of Theorem 2

Without loss of generality we can assume that the game is not trivial and that  $\mathcal{X}(z)$  is finite for any  $z \in \mathcal{Z}$  since otherwise if  $\mathcal{X}(z)$  is a polyhedron, the maximum in (2) must be attained at an extreme point of  $\mathcal{X}(z)$  ( $\ell$  is bounded by Assumption 1) and there are finitely many such points for any  $z$ . Moreover, we can also assume that  $\mathcal{Z}$  is discrete by Lemma 4 since, borrowing the notations of Lemma 4, we have:

$$R_T(\ell, \mathcal{Z}, \mathcal{F}) \geq R_T(\ell, \tilde{\mathcal{Z}}, \mathcal{F}).$$

Write  $\mathcal{Z} = \{z_n \mid n \leq N\}$  and denote by  $p_0$  the uniform distribution on  $\mathcal{Z}$ , i.e.  $p_0 = \frac{1}{N} \sum_{n=1}^N \delta_{z_n}$ , where  $\delta_{z_n}$  is the Dirac distribution supported at  $z_n$ . We may assume that there is a single equivalence class in  $\argmin_{f \in \mathcal{F}} \mathbb{E}_{p_0}[\ell(Z, f)]$ , otherwise we are done by Lemma 2. Take  $f^* \in \argmin_{f \in \mathcal{F}} \mathbb{E}_{p_0}[\ell(Z, f)]$ . Since the game  $(\ell, \mathcal{Z}, \mathcal{F})$  is not trivial, there exists  $z_k$  in  $\mathcal{Z}$  and  $f^{**}$  in  $\mathcal{F}$  such that  $\ell(z_k, f^{**}) < \ell(z_k, f^*)$ . Therefore, we can find  $\epsilon > 0$  small enough such that  $(N-1)\epsilon < 1$  and:

$$(1 - (N-1)\epsilon)\ell(z_k, f^{**}) + \sum_{n \neq k} \epsilon \ell(z_n, f^{**}) < (1 - (N-1)\epsilon)\ell(z_k, f^*) + \sum_{n \neq k} \epsilon \ell(z_n, f^*).$$

Define  $p_1$  as the corresponding distribution, i.e.  $p_1 = (1 - (N-1)\epsilon)\delta_{z_k} + \epsilon \sum_{n \neq k} \delta_{z_n}$ . By construction, the equivalence class of  $f^*$  is not in  $\argmin_{f \in \mathcal{F}} \mathbb{E}_{p_1}[\ell(Z, f)]$ . Once again, without loss of generality, we may assume that there is a single equivalence class in  $\argmin_{f \in \mathcal{F}} \mathbb{E}_{p_1}[\ell(Z, f)]$ , otherwise we are done by Lemma 2. Moreover, we can now redefine  $f^{**}$  as a representative of the only equivalence class contained in  $\argmin_{f \in \mathcal{F}} \mathbb{E}_{p_1}[\ell(Z, f)]$ . We now move on to show that there must exist  $\alpha \in (0, 1)$  such that there are at least two equivalence classes in  $\argmin_{f \in \mathcal{F}} \mathbb{E}_{p_\alpha}[\ell(Z, f)]$ ,

where the distribution  $p_\alpha$  is defined as  $p_\alpha = (1 - \alpha)p_0 + \alpha p_1$ . Observe that  $\min_{f \in \mathcal{F}} \mathbb{E}_{p_\alpha}[\ell(Z, f)]$  can be written as the linear program:

$$\begin{aligned} \min_{q_1, \dots, q_N, f} \quad & q \cdot ((1 - \alpha)x_0 + \alpha x_1) \\ \text{subject to} \quad & q = (q_1, \dots, q_N) \\ & q_n \geq (C(z_n)f + c(z_n))^\top x \quad \forall x \in \mathcal{X}(z_n), \forall n = 1, \dots, N \\ & f \in \mathcal{F}, q_1, \dots, q_N \in \mathbb{R} \end{aligned} \quad (3)$$

where  $x_0$  and  $x_1$  are vectors of size  $N$  defined as follows:

$$x_0 = \frac{1}{N}(1, \dots, 1)$$

and

$$x_1 = (1 - (N - 1)\epsilon)(0, \dots, 0, 1, 0, \dots) + \epsilon(1, \dots, 1, 0, 1, \dots, 1),$$

i.e. all the components of  $x_1$  are equal to  $\epsilon$  except for the  $k$ th one which is equal to  $(1 - (N - 1)\epsilon)$ . We are interested in the function  $\phi : \alpha \rightarrow \arg\min_{f \in \mathcal{F}} \mathbb{E}_{p_\alpha}[\ell(Z, f)]$ . For any  $f \in \mathcal{F}$ , define  $I(f) = \{\alpha \in [0, 1] \mid f \in \phi(\alpha)\}$ . Since  $\alpha \rightarrow \mathbb{E}_{p_\alpha}[\ell(Z, f)]$  is linear in  $\alpha$ ,  $I(f) = \{\alpha \in [0, 1] \mid f \in \phi(\alpha)\}$  is a closed interval in  $[0, 1]$  for any  $f \in \mathcal{F}$ . Moreover,  $\mathcal{F}$  being a polyhedron, the feasible set of (3) is also a polyhedron, hence it has a finitely many extreme points. We denote by  $\{f_1, \dots, f_L\}$  the projection of the set of extreme points onto the  $f$  coordinate. Since (3) is a linear program, this shows that, for any  $\alpha \in [0, 1]$ , there exists  $l \in \{1, \dots, L\}$  such that  $f_l \in \phi(\alpha)$ . Therefore, we can write  $[0, 1] = \cup_{l=1}^L I(f_l)$ . We can further simplify this description by assuming that the  $f_l$ 's belong to different equivalence classes (because  $I(f) = I(f')$  if  $f$  is equivalent to  $f'$ ). Now observe that if  $I(f_l) \cap I(f_j) \neq \emptyset$  for all any  $l \neq j \leq K$ , then there are two classes of equivalence in  $\arg\min_{f \in \mathcal{F}} \mathbb{E}_{p_\alpha}[\ell(Z, f)]$  for any  $\alpha \in I(f_l) \cap I(f_j)$  and we are done. Suppose by contradiction that we cannot find such a pair of indices. Because the only way to partition  $[0, 1]$  into  $L < \infty$  non-overlapping closed intervals is to have  $L = 1$ , we get  $[0, 1] = I(f_1)$ . This implies that  $f^*$  and  $f^{**}$  belong to the same equivalence class, a contradiction.

#### 4.7 Alternative proof of Theorem 2 by exhibiting an equalizing strategy when $\ell$ is linear

Using Lemma 1, we can assume without loss of generality that  $\mathcal{Z}$  is convex. When  $\ell$  is linear, the procedure developed in the proof of Theorem 2 boils down to finding a point  $z \in \text{int}(\mathcal{Z})$  such that  $|\arg\min_{f \in \mathcal{F}} z^\top f| > 1$  and, with further examination, we can also guarantee that there exists  $\epsilon > 0, e \in \mathbb{R}^n$  and  $f_1, f_2 \in \arg\min_{f \in \mathcal{F}} z^\top f$  such that  $f_1 \in \arg\min_{f \in \mathcal{F}} (z - xe)^\top f$  while  $f_2 \notin \arg\min_{f \in \mathcal{F}} (z - xe)^\top f$  for all  $x \in (0, \epsilon]$  and symmetrically for  $x \in [-\epsilon, 0)$ . Consider a randomized opponent  $Z_t = z + (\epsilon_t e)$  for  $(\epsilon_t)_{t=1, \dots, T}$  i.i.d. Rademacher random variables. Then for any player's strategy:

$$\mathbb{E}[r_T((Z_t)_{t=1, \dots, T}, (f_t)_{t=1, \dots, T})] = \sum_{t=1}^T \mathbb{E}[Z_t]^\top f_t - \mathbb{E}[\inf_{f \in \mathcal{F}} f^\top \sum_{t=1}^T Z_t].$$

This yields:

$$\mathbb{E}[r_T((Z_t)_{t=1, \dots, T}, (f_t)_{t=1, \dots, T})] = \sum_{t=1}^T z^\top f_t - \mathbb{E}[\inf_{f \in \mathcal{F}} f^\top \sum_{t=1}^T Z_t].$$

We can lower bound the last quantity by:

$$\mathbb{E}[r_T((Z_t)_{t=1, \dots, T}, (f_t)_{t=1, \dots, T})] \geq T(z^\top f_1) - T\mathbb{E}[\inf_{f \in \mathcal{F}} f^\top (z + (\epsilon \cdot \frac{\sum_{t=1}^T \epsilon_t}{T})e)],$$

as  $f_1 \in \arg\min_{f \in \mathcal{F}} z^\top f$ , but we could have equivalently picked  $f_2$  as  $f_1^\top z = f_2^\top z$ . Furthermore, as  $|\frac{\sum_{t=1}^T \epsilon_t}{T}| \leq 1$ ,  $f_1$  is optimal in the inner optimization problem when  $\sum_{t=1}^T \epsilon_t \leq 0$  while  $f_2$  is optimal when  $\sum_{t=1}^T \epsilon_t \geq 0$ . Hence:

$$\begin{aligned} \mathbb{E}[r_T((Z_t)_{t=1, \dots, T}, (f_t)_{t=1, \dots, T})] &\geq T(z^\top f_1) - T\mathbb{E}[f_1^\top (z + (\epsilon \cdot \frac{\sum_{t=1}^T \epsilon_t}{T})e) \cdot 1_{\sum_{t=1}^T \epsilon_t \leq 0} + \\ &\quad f_2^\top (z + (\epsilon \cdot \frac{\sum_{t=1}^T \epsilon_t}{T})e) \cdot 1_{\sum_{t=1}^T \epsilon_t \geq 0}]. \end{aligned}$$

Observe that the term  $T(z^\top f_1)$  cancels out and we get:

$$\mathbb{E}[r_T((Z_t)_{t=1,\dots,T}, (f_t)_{t=1,\dots,T})] \geq \frac{\mathbb{E}[\|\sum_{t=1}^T \epsilon_t\|]}{T} \cdot \epsilon \cdot (f_1^\top e - f_2^\top e).$$

By Khintchine's inequality  $\mathbb{E}[\|\sum_{t=1}^T \epsilon_t\|] \geq \frac{1}{\sqrt{2}}\sqrt{T}$ . Moreover  $f_1^\top e - f_2^\top e > 0$  because  $f_2 \in \operatorname{argmin}_{f \in \mathcal{F}} (z + \epsilon e)^\top f$  while  $f_1$  does not and  $f_1^\top z = f_2^\top z$ . We finally derive

$$\mathbb{E}[r_T((Z_t)_{t=1,\dots,T}, (f_t)_{t=1,\dots,T})] \geq \frac{(f_1^\top e - f_2^\top e)}{\sqrt{2}}\sqrt{T}.$$

This enables us to conclude  $R_T = \Omega(\sqrt{T})$  as this shows that for any player's strategy, there exists a sequence  $z_1, \dots, z_T$  such that

$$r_T((z_t)_{t=1,\dots,T}, (f_t)_{t=1,\dots,T}) \geq \mathbb{E}[r_T((Z_t)_{t=1,\dots,T}, (f_t)_{t=1,\dots,T})].$$

#### 4.8 Proof of Theorem 3

Straightforward from Theorem 2 since  $\ell$  is jointly continuous.

#### 4.9 Proof of Lemma 5

Using Lemma 1, we can assume that  $\mathcal{Z}$  is convex. Consider  $f_1 \neq f_2 \in \mathcal{F}$  and define  $e = \frac{f_1 - f_2}{\|f_1 - f_2\|}$ . Since  $0 \in \operatorname{int}(\mathcal{Z})$ , there exists  $\epsilon > 0$  such that  $\epsilon e$  and  $-\epsilon e$  are in  $\mathcal{Z}$ . We restrict the opponent's decision set by imposing that, at any round  $t$ , the opponent's move be  $y_t \epsilon e$  for  $y_t \in \tilde{\mathcal{Z}} = \{-1, 1\}$ . Since  $\ell(y_t \epsilon e, f)$  only depends on  $f$  through the scalar product between  $f$  and  $e$ , the player's decision set can equivalently be described by  $\tilde{\mathcal{F}} = \{f^\top e \mid f \in \mathcal{F}\}$  which is a closed interval (since  $\mathcal{F}$  is convex and compact) and thus a polyhedron. Defining a new loss function as  $\tilde{\ell}(y, f) = y \epsilon f$ , we have:

$$R_T(\ell, \mathcal{Z}, \mathcal{F}) \geq R_T(\tilde{\ell}, \tilde{\mathcal{Z}}, \tilde{\mathcal{F}}).$$

Observe that the game  $(\tilde{\ell}, \tilde{\mathcal{Z}}, \tilde{\mathcal{F}})$  is linear and not trivial, otherwise there would exist  $f^*$  such that  $e^\top f^* \leq e^\top f_2$  and  $-e^\top f^* \leq -e^\top f_1$  which would imply  $\|e\| = 0$ . With Theorem 3, we conclude  $R_T(\tilde{\ell}, \tilde{\mathcal{Z}}, \tilde{\mathcal{F}}) = \Omega(\sqrt{T})$  and thus  $R_T(\ell, \mathcal{Z}, \mathcal{F}) = \Omega(\sqrt{T})$ .

#### 4.10 Proof of Theorem 4

A common inequality on the regret incurred for the FTL strategy is:

$$r_T((z_t)_{t=1,\dots,T}, (f_t)_{t=1,\dots,T}) \leq \sum_{t=1}^T z_t^\top (f_t - f_{t+1}),$$

We use sensitivity analysis to control this last quantity. Specifically we show that the mapping  $\phi : z \rightarrow \operatorname{argmin}_{f \in \mathcal{F}(f)=0} z^\top f$  is Lipschitz on  $\mathcal{Z}$ . Using this property:

$$\begin{aligned} r_T((z_t)_{t=1,\dots,T}, (f_t)_{t=1,\dots,T}) &\leq \sum_{t=1}^T \|z_t\| \|f_t - f_{t+1}\| \\ &= O\left(\sum_{t=1}^T \left\| \frac{1}{t-1} \sum_{\tau=1}^{t-1} z_\tau - \frac{1}{t} \sum_{\tau=1}^t z_\tau \right\|\right) \\ &= O\left(\sum_{t=1}^T \left\| \frac{1}{t(t-1)} \sum_{\tau=1}^{t-1} z_\tau - \frac{1}{t} z_t \right\|\right) \\ &= O\left(\sum_{t=1}^T \frac{1}{t(t-1)} \left\| \sum_{\tau=1}^{t-1} z_\tau \right\| + \frac{1}{t} \|z_t\|\right) \\ &= O\left(\sum_{t=1}^T \frac{1}{t}\right) \\ &= O(\log(T)), \end{aligned}$$

since  $\mathcal{Z}$  is compact.

We now move on to show that  $\phi$  is Lipschitz. As  $\text{conv}(\mathcal{Z})$  is closed and convex, we can strictly separate 0 from  $\text{conv}(\mathcal{Z})$ . Hence, there exists  $a \neq 0 \in \mathbb{R}^d$  and  $c > 0$  such that  $a \cdot z > c, \forall z \in \mathcal{Z}$ . We get  $\|z\| \geq \frac{c}{\|a\|} > 0 \forall z \in \mathcal{Z}$ . Let us use the shorthand  $C = \frac{c}{\|a\|}$ . Take  $(z_1, z_2) \in \mathcal{Z}^2$  and  $(f(z_1), f(z_2)) \in \phi(z_1) \times \phi(z_2)$ . Observe that the constraint qualifications are automatically satisfied at  $f(z_1)$  and  $f(z_2)$  as  $\nabla F$  cannot vanish on  $\{f \mid F(f) = 0\}$  since  $F$  does not attain its minimum on this set ( $\mathcal{F}$  is assumed to contain at least two points). Hence, there exist  $\lambda_1, \lambda_2 \geq 0$  such that  $z_1 + \lambda_1 \nabla F(f(z_1)) = 0$  and  $z_2 + \lambda_2 \nabla F(f(z_2)) = 0$ . As  $z_1, z_2 \neq 0$ , we must have  $\lambda_1, \lambda_2 \neq 0$ . We obtain  $\nabla F(f_1) = -\frac{1}{\lambda_1} z_1$  and  $\nabla F(f(z_2)) = -\frac{1}{\lambda_2} z_2$ . Since  $F$  is strongly convex, there exists  $\beta > 0$  such that:

$$(\nabla F(f') - \nabla F(f''))^\top (f' - f'') \geq \beta \|f' - f''\|^2,$$

for all  $f', f'' \in \mathcal{F}$ . Applying this inequality for  $f' = f(z_1)$  and  $f'' = f(z_2)$ , we obtain:

$$\left(\frac{1}{\lambda_2} z_2 - \frac{1}{\lambda_1} z_1\right)^\top (f(z_1) - f(z_2)) \geq \beta \|f(z_1) - f(z_2)\|^2.$$

We can break down the last expression in two pieces:

$$\frac{1}{\lambda_2} z_2^\top (f(z_1) - f(z_2)) + \frac{1}{\lambda_1} z_1^\top (f(z_2) - f(z_1)) \geq \beta \|f(z_1) - f(z_2)\|^2.$$

Observe that  $z_2^\top (f(z_1) - f(z_2)) \geq 0$  since both  $f(z_1)$  and  $f(z_2)$  belong to  $\mathcal{F}$  and since  $f(z_2)$  is the minimizer of  $z_2^\top f$  for  $f$  ranging in  $\mathcal{F}$ . Symmetrically,  $z_1^\top (f(z_2) - f(z_1)) \geq 0$ . Note that  $\frac{1}{\lambda_1} = \frac{1}{\|z_1\|} = \frac{\|\nabla F(f(z_1))\|}{\|z_1\|}$ . As  $\nabla F$  is continuous and  $\mathcal{F}$  is compact, there exists  $K > 0$  such that  $\|\nabla F(f)\| \leq K$  for any  $f \in \mathcal{F}$ . Hence, we get  $\frac{1}{\lambda_1} \leq \frac{K}{C}$  and the same inequality holds for  $\lambda_2$ . Plugging this upper bound back into the last inequality yields:

$$\frac{K}{C} (z_2 - z_1)^\top (f(z_1) - f(z_2)) \geq \beta \|f(z_1) - f(z_2)\|^2.$$

Using the Cauchy-Schwarz inequality and simplifying on both sides by  $\|f(z_2) - f(z_1)\|$  yields:

$$\frac{K}{\beta C} \|z_2 - z_1\| \geq \|f(z_1) - f(z_2)\|,$$

i.e.  $\frac{K}{\beta C} \|z_2 - z_1\| \geq \|\phi(z_1) - \phi(z_2)\|$ .

#### 4.11 Proof of Lemma 6

Using Lemma 1, we can assume that  $\mathcal{Z}$  is convex. Since  $\mathcal{Z}$  is compact and convex and since  $0 \notin \mathcal{Z}$ , we can strictly separate 0 from  $\mathcal{Z}$  and find  $z^* \neq 0$  such that  $\mathcal{Z} \subseteq B_2(z^*, \alpha \|z^*\|)$  with  $\alpha < 1$ . By rescaling  $\mathcal{Z}$ , we can assume that  $\alpha \|z^*\| = 1$  and  $\|z^*\| > 1$ . In the sequel,  $\sigma_{t-1}$  serves as a shorthand for  $\sigma(Z_1, \dots, Z_{t-1})$ . We prove more generally that, for any choice of random variables  $(Z_1, \dots, Z_T)$  such that  $\mathbb{E}[Z_t | \sigma_{t-1}]$  is constant almost surely, the lower bound on regret derived from Theorem 1 is  $O(1)$ . Write  $Z_t = z^* + W_t$  and  $\mathbb{E}[W_t | \sigma_{t-1}] = c_t$  with  $\|W_t\| \leq 1$  and  $\|c_t\| \leq 1$ . Define  $w^* = T \cdot z^* + \sum_{t=1}^T c_t$ . Observe that  $\|w^*\| \geq T \cdot \|z^*\| - \sum_{t=1}^T \|c_t\| \geq T \cdot (\|z^*\| - 1) > 0$ . Write  $W_t = X_t \frac{w^*}{\|w^*\|} + \tilde{W}_t + c_t$  with  $\tilde{W}_t^\top w^* = 0$ . Projecting down the equality  $\mathbb{E}[W_t - c_t | \sigma_{t-1}] = 0$  onto  $w^*$ , we get  $\mathbb{E}[X_t | \sigma_{t-1}] = 0$  and  $\mathbb{E}[\tilde{W}_t | \sigma_{t-1}] = 0$ . The bound that results from an application of Theorem 1 is:

$$R_T \geq \mathbb{E}[\|w^* + \sum_{t=1}^T W_t - c_t\|] - \sum_{t=1}^T \|z^* + c_t\|.$$

We now focus on finding an upper bound on the right-hand side. Expanding the first term yields:

$$\|w^* + \sum_{t=1}^T W_t - c_t\| = \sqrt{\left(1 + \sum_{t=1}^T \frac{X_t}{\|w^*\|}\right)^2 \cdot \|w^*\|^2 + \left\|\sum_{t=1}^T \tilde{W}_t\right\|^2}.$$



By concavity of the squared root function:

$$\mathbb{E}[\|w^* + \sum_{t=1}^T W_t - c_t\|] \leq \sqrt{\|w^*\|^2 \cdot \mathbb{E}[(1 + \sum_{t=1}^T \frac{X_t}{\|w^*\|})^2] + \mathbb{E}[\|\sum_{t=1}^T \tilde{W}_t\|^2]}.$$

We expand the two inner terms:

$$\mathbb{E}[(1 + \sum_{t=1}^T \frac{X_t}{\|w^*\|})^2] = 1 + 2 \sum_{t=1}^T \frac{\mathbb{E}[X_t]}{\|w^*\|} + \frac{1}{\|w^*\|^2} \cdot \mathbb{E}[(\sum_{t=1}^T X_t)^2].$$

Looking at each term individually, we have  $\mathbb{E}[X_t] = \mathbb{E}[\mathbb{E}[X_t|\sigma_{t-1}]] = 0$  and  $\mathbb{E}[(\sum_{t=1}^T X_t)^2] = \mathbb{E}[(\sum_{t=1}^{T-1} X_t)^2] + 2\mathbb{E}[X_T \cdot (\sum_{t=1}^{T-1} X_t)] + \mathbb{E}[X_T^2]$ , yet  $\mathbb{E}[X_T \cdot (\sum_{t=1}^{T-1} X_t)] = \mathbb{E}[\mathbb{E}[X_T|\sigma_{T-1}] \cdot (\sum_{t=1}^{T-1} X_t)] = 0$ . Hence,  $\mathbb{E}[(1 + \sum_{t=1}^T \frac{X_t}{\|w^*\|})^2] = 1 + \frac{\mathbb{E}[\sum_{t=1}^T X_t^2]}{\|w^*\|^2}$ . Similarly  $\mathbb{E}[\|\sum_{t=1}^T \tilde{W}_t\|^2] = \sum_{t=1}^T \mathbb{E}[\|\tilde{W}_t\|^2]$ . We obtain:

$$\mathbb{E}[\|w^* + \sum_{t=1}^T W_t - c_t\|] \leq \sqrt{\|w^*\|^2 + \sum_{t=1}^T \mathbb{E}[X_t^2 + \|\tilde{W}_t\|^2]}.$$

Remark that  $\|W_t - c_t\| \leq \|W_t\| + \|c_t\| \leq 2$ . Hence,  $X_t^2 + \|\tilde{W}_t\|^2 \leq 2$ . We obtain:

$$\mathbb{E}[\|w^* + \sum_{t=1}^T W_t - c_t\|] \leq \sqrt{\|w^*\|^2 + 2T}.$$

We have  $\sqrt{\|w^*\|^2 + 2T} = \|w^*\| \cdot \sqrt{1 + \frac{2T}{\|w^*\|^2}} \leq \|w^*\| + \frac{T}{\|w^*\|}$  for  $T$  big enough as  $\|w^*\| \geq T \cdot (\|z^*\| - 1)$ . Yet  $\|w^*\| = \|\sum_{t=1}^T z^* + c_t\| \leq \sum_{t=1}^T \|z^* + c_t\|$ . Hence, the lower bound derived is:

$$\mathbb{E}[\|w^* + \sum_{t=1}^T W_t - c_t\|] - \sum_{t=1}^T \|z^* + c_t\| \leq \frac{T}{\|w^*\|} \leq \frac{1}{\|z^*\| - 1} = O(1).$$

#### 4.12 Proof of Theorem 5

Since  $\mathcal{Z}$  has non-empty interior, we can find  $z^* \neq 0$  and  $\alpha \in (0, \frac{1}{32}]$  such that  $B_2(z^*, \alpha \|z^*\|) \subseteq \mathcal{Z}$ . Define  $e$  as a unit vector orthogonal to  $z^*$  and  $\tilde{\mathcal{Z}} = \{z \mid z = z^* + (w\alpha \|z^*\|)e, |w| \leq 1\}$ . Since  $\tilde{\mathcal{Z}} \subseteq \mathcal{Z}$ , we have:

$$R_T(\ell, \mathcal{Z}, \mathcal{F}) \geq R_T(\ell, \tilde{\mathcal{Z}}, \mathcal{F}),$$

and we can focus on developing a  $\Omega(\log(T))$  lower bound on regret for the game  $(\ell, \tilde{\mathcal{Z}}, \mathcal{F})$ . Using the minimax reformulation of Theorem 1, we have:

$$\begin{aligned} R_T(\ell, \tilde{\mathcal{Z}}, \mathcal{F}) &= \sup_p \mathbb{E} \left[ - \sum_{t=1}^T \|z^* + (\alpha \|z^*\| \mathbb{E}[W_t|W_1, \dots, W_{t-1}])e\| + \left\| Tz^* + (\alpha \|z^*\| \sum_{t=1}^T W_t)e \right\| \right] \\ &= \sup_p \mathbb{E} \left[ - \sum_{t=1}^T \sqrt{\|z^*\|^2 + (\alpha \|z^*\| \mathbb{E}[W_t|W_1, \dots, W_{t-1}])^2} + \sqrt{T^2 \|z^*\|^2 + (\alpha \|z^*\| \sum_{t=1}^T W_t)^2} \right] \end{aligned}$$

where the supremum is taken over the distribution  $p$  of the random variable  $(W_1, \dots, W_T)$  in  $[-1, 1]^T$ . Rearranging this expression yields:

$$\begin{aligned}
R_T(\ell, \tilde{Z}, \mathcal{F}) &= \|z^*\| \sup_p \mathbb{E} \left[ T \sqrt{1 + \left( \alpha \frac{\sum_{t=1}^T W_t}{T} \right)^2} - \sum_{t=1}^T \sqrt{1 + (\alpha \mathbb{E}[W_t | W_1, \dots, W_{t-1}])^2} \right] \\
&= \|z^*\| \sup_p \{ \mathbb{E} [ T(1 + \sum_{n=1}^{\infty} \binom{\frac{1}{2}}{n} \alpha^{2n} (\frac{\sum_{t=1}^T W_t}{T})^{2n}) \\
&\quad - \sum_{t=1}^T (1 + \sum_{n=1}^{\infty} \binom{\frac{1}{2}}{n} \alpha^{2n} \mathbb{E}[W_t | W_1, \dots, W_{t-1}]^{2n}) ] \} \\
&= \|z^*\| \sup_p \{ \frac{\alpha^2}{2} \mathbb{E} \left[ \frac{(\sum_{t=1}^T W_t)^2}{T} - \sum_{t=1}^T \mathbb{E}[W_t | W_1, \dots, W_{t-1}]^2 \right] \\
&\quad + \sum_{n=2}^{\infty} \binom{\frac{1}{2}}{n} \alpha^{2n} \mathbb{E} \left[ \frac{(\sum_{t=1}^T W_t)^{2n}}{T^{2n-1}} - \sum_{t=1}^T \mathbb{E}[W_t | W_1, \dots, W_{t-1}]^{2n} \right] \},
\end{aligned}$$

where the second equality results from a series expansion (valid since  $(\alpha \mathbb{E}[W_t | W_1, \dots, W_{t-1}])^2, (\alpha \frac{\sum_{t=1}^T W_t}{T})^2 \leq \alpha^2 < 1$ ) and the third inequality is derived from Fubini, observing that:

$$\sum_{n=1}^{\infty} \left| \binom{\frac{1}{2}}{n} \right| \alpha^{2n} \mathbb{E} \left[ \left( \frac{\sum_{t=1}^T W_t}{T} \right)^{2n} \right] \leq \sum_{n=1}^{\infty} \alpha^{2n} = \frac{1}{1 - \alpha^2} < \infty$$

and similarly:

$$\sum_{n=1}^{\infty} \left| \binom{\frac{1}{2}}{n} \right| \alpha^{2n} \mathbb{E} [\mathbb{E}[W_t | W_1, \dots, W_{t-1}]^{2n}] \leq \sum_{n=1}^{\infty} \alpha^{2n} = \frac{1}{1 - \alpha^2} < \infty.$$

Interestingly, the first-order term of this series expansion, i.e.

$$\mathbb{E} \left[ \frac{(\sum_{t=1}^T W_t)^2}{T} - \sum_{t=1}^T \mathbb{E}[W_t | W_1, \dots, W_{t-1}]^2 \right],$$

is precisely the expression of the minimax regret for the game  $(\ell(z, f) = (z - f)^2, [-1, 1], [-1, 1])$  which is known to have optimal regret  $\Theta(\log(T))$ , see Section 7.3 of [2]. This motivates the introduction of the probability distribution  $p$  used in [2] to establish the  $\Omega(\log(T))$  lower bound. Specifically, we use the conditional distributions:

$$p_t(W_t = w | W_1, \dots, W_{t-1}) = \begin{cases} \frac{1 + c_t W_{1:t-1}}{2} & \text{if } w = 1 \\ \frac{1 - c_t W_{1:t-1}}{2} & \text{if } w = -1 \end{cases} \quad t = 2, \dots, T$$

where  $W_{1:t-1} = \sum_{\tau=1}^{t-1} W_\tau$  and the sequence  $(c_t)_{t=1, \dots, T}$  is recursively defined as:

$$\begin{aligned}
c_T &= \frac{1}{T} \\
c_{t-1} &= c_t + c_t^2 \quad t = T, \dots, 2.
\end{aligned}$$

Together with  $W_1$  taken as a Rademacher random variable, these conditional distributions define a joint distribution  $p$  as it can be shown that  $c_t \in [0, \frac{1}{t}]$ . Abernethy et al. [2] show that:

$$\mathbb{E} \left[ \frac{(\sum_{t=1}^T W_t)^2}{T} - \sum_{t=1}^T \mathbb{E}[W_t | W_1, \dots, W_{t-1}]^2 \right] = \log(T) + O(\log \log(T)). \quad (4)$$

Hence, it remains to control the terms of order  $n \geq 2$  in the series expansion. First observe that, by definition:

$$\begin{aligned}
\mathbb{E}[W_{1:T}^{2n}] &= \mathbb{E}[\mathbb{E}[W_{1:T}^{2n} | W_1, \dots, W_{T-1}]] \\
&= \mathbb{E} \left[ \frac{1 + c_T W_{1:T-1}}{2} (W_{1:T-1} + 1)^{2n} + \frac{1 - c_T W_{1:T-1}}{2} (W_{1:T-1} - 1)^{2n} \right] \\
&= 1 + \sum_{k=1}^n \binom{2n}{2k} c_T \mathbb{E}[(W_{1:T-1})^{2k}],
\end{aligned}$$

which implies that:

$$\begin{aligned} |c_T^{2n-1}\mathbb{E}[W_{1:T}^{2n}] - (c_T^{2n-1} + 2nc_T^{2n})\mathbb{E}[W_{1:T-1}^{2n}]| &\leq \binom{2n}{2(n-1)}c_T + 2\sum_{k=0}^n \binom{2n}{2k}c_T^2 \\ &\leq 2n^2c_T + 24^n c_T^2, \end{aligned} \quad (5)$$

since  $c_T|W_{1:T-1}| \leq 1$ . Additionally, we have, using the recursive definition of the sequence  $(c_t)_{t=1,\dots,T}$ :

$$\begin{aligned} c_{T-1}^{2n-1} &= (c_T + c_T^2)^{2n-1} \\ &= \sum_{k=0}^{2n-1} \binom{2n-1}{k} c_T^{2n-1+k}, \end{aligned}$$

which implies:

$$|c_{T-1}^{2n-1} - (c_T^{2n-1} + (2n-1)c_T^{2n})| \leq 2n^2c_T^{2n+1} + 4^n c_T^{2n+2}. \quad (6)$$

Using  $\mathbb{E}[W_T|W_1, \dots, W_{T-1}] = c_T W_{1:T-1}$ , we get:

$$\begin{aligned} &|\mathbb{E}[c_T^{2n-1}W_{1:T}^{2n} - \sum_{t=1}^T \mathbb{E}[W_t|W_1, \dots, W_{t-1}]^{2n}]| \\ &\leq |\mathbb{E}[(c_T^{2n-1} + 2(n-1)c_T^{2n})W_{1:T-1}^{2n}] - \sum_{t=1}^{T-1} \mathbb{E}[W_t|W_1, \dots, W_{t-1}]^{2n}]| \\ &\quad + 2n^2c_T + 24^n c_T^2 \\ &\leq |\mathbb{E}[c_{T-1}^{2n-1}W_{1:T-1}^{2n} - \sum_{t=1}^{T-1} \mathbb{E}[W_t|W_1, \dots, W_{t-1}]^{2n}]| \\ &\quad + (2n^2c_T^{2n+1} + 4^n c_T^{2n+2})\mathbb{E}[W_{1:T-1}^{2n}] + 2n^2c_T + 24^n c_T^2 \\ &\leq 4n^2c_T + 34^n c_T^2 \\ &\leq 4n^2\frac{1}{T} + 34^n\frac{1}{T^2}, \end{aligned}$$

where the first (resp. second) inequality is obtained by applying (5) (resp. (6)) and the fifth inequality is derived using  $c_T \in [0, \frac{1}{T}]$  and  $|W_{1:T-1}| \leq T-1$ . By induction on  $t$ , we get:

$$\begin{aligned} \left| \mathbb{E} \left[ c_T^{2n-1}W_{1:T}^{2n} - \sum_{t=1}^T \mathbb{E}[W_t|W_1, \dots, W_{t-1}]^{2n} \right] \right| &\leq 4n^2 \sum_{t=1}^T \frac{1}{t} + 34^n \sum_{t=1}^T \frac{1}{t^2} \\ &\leq 4n^2 \log(T) + 4^n \frac{\pi^2}{2}. \end{aligned}$$

Bringing everything together, we derive:

$$\begin{aligned}
& \left| \sum_{n=2}^{\infty} \binom{\frac{1}{2}}{n} \alpha^{2n} \mathbb{E} \left[ \frac{(\sum_{t=1}^T W_t)^{2n}}{T^{2n-1}} - \sum_{t=1}^T \mathbb{E}[W_t | W_1, \dots, W_{t-1}]^{2n} \right] \right| \\
& \leq 4 \left( \sum_{n=2}^{\infty} \binom{\frac{1}{2}}{n} \alpha^{2n} n^2 \right) \log(T) \\
& \quad + \left( \sum_{n=2}^{\infty} \binom{\frac{1}{2}}{n} (2\alpha)^{2n} \right) \frac{\pi^2}{2} \\
& \leq 8\alpha^4 \left( \sum_{n=2}^{\infty} n(n-1)(\alpha^2)^{n-2} \right) \log(T) \\
& \quad + \left( \sum_{n=0}^{\infty} (2\alpha)^{2n} \right) \frac{\pi^2}{2} \\
& \leq 8 \frac{\alpha^4}{(1-\alpha^2)^3} \log(T) + \frac{\pi^2}{2(1-2\alpha)} \\
& \leq 8 \frac{\alpha^4}{(1-\alpha^2)^3} \log(T) + \pi^2,
\end{aligned}$$

since  $\alpha \leq \frac{1}{4}$ . Using (4), we conclude that:

$$R_T(\ell, \tilde{\mathcal{Z}}, \mathcal{F}) \geq \frac{\|z^*\| \alpha^2}{2} (1 - 16 \frac{\alpha^2}{(1-\alpha^2)^3}) \log(T) + O(\log \log(T)),$$

which implies that  $R_T(\ell, \tilde{\mathcal{Z}}, \mathcal{F}) = \Omega(\log(T))$  as  $\alpha^2(1 - 16 \frac{\alpha^2}{(1-\alpha^2)^3}) > 0$  for  $\alpha \in (0, \frac{1}{32}]$ .

#### 4.13 Proof of Lemma 7

The proof is along the same lines as for Theorem 4. We start with the same inequality:

$$r_T((z_t)_{t=1, \dots, T}, (f_t)_{t=1, \dots, T}) \leq \sum_{t=1}^T z_t^\top (f_t - f_{t+1}),$$

and use sensitivity analysis to control this last quantity. Specifically, we show that the mapping  $\phi : z \rightarrow \operatorname{argmin}_{f \in \mathcal{F}} z^\top f$  is  $\frac{1}{q-1}$ -Hölder continuous on  $\mathcal{Z}$ , i.e. there exists  $c > 0$  such that:

$$\|\phi(z_1) - \phi(z_2)\|_2 \leq c \|z_1 - z_2\|_2^{\frac{1}{q-1}} \quad \forall (z_1, z_2) \in \mathcal{Z}^2.$$

Using this property, we get:

$$\begin{aligned}
r_T((z_t)_{t=1, \dots, T}, (f_t)_{t=1, \dots, T}) & \leq \sum_{t=1}^T \|z_t\|_2 \|f_t - f_{t+1}\|_2 \\
& = O\left(\sum_{t=1}^T \left\| \frac{1}{t-1} \sum_{\tau=1}^{t-1} z_\tau - \frac{1}{t} \sum_{\tau=1}^t z_\tau \right\|_2^{\frac{1}{q-1}}\right) \\
& = O\left(\sum_{t=1}^T \left\| \frac{1}{t(t-1)} \sum_{\tau=1}^{t-1} z_\tau - \frac{1}{t} z_t \right\|_2^{\frac{1}{q-1}}\right) \\
& = O\left(\sum_{t=1}^T \left( \frac{1}{t(t-1)} \left\| \sum_{\tau=1}^{t-1} z_\tau \right\|_2 + \frac{1}{t} \|z_t\|_2 \right)^{\frac{1}{q-1}}\right) \\
& = O\left(\sum_{t=1}^T \frac{1}{t^{\frac{1}{q-1}}}\right),
\end{aligned}$$

from which we derive that  $r_T((z_t)_{t=1,\dots,T}, (f_t)_{t=1,\dots,T}) = O(\log(T))$  if  $q = 2$  and  $r_T((z_t)_{t=1,\dots,T}, (f_t)_{t=1,\dots,T}) = O(T^{\frac{q-2}{q-1}})$  if  $q \in (2, 3]$ .

We now move on to show that  $\phi$  is  $\frac{1}{q-1}$ -Hölder continuous. Just like in Theorem 4, we can find  $A > 0$  such that  $\|z\|_2 \geq A$  for all  $z \in \mathcal{Z}$ . Take  $(z_1, z_2) \in \mathcal{Z}^2$  and  $(f_1, f_2) \in \phi(z_1) \times \phi(z_2)$ . Since we are optimizing a linear function, we may assume, without loss of generality, that  $C = 1$  and  $f_1$  and  $f_2$  lie on the boundary of  $\mathcal{F}$ , i.e.  $\|f_1\|_{\mathcal{F}} = \|f_2\|_{\mathcal{F}} = 1$ . By definition, we have:

$$\left\| \frac{f_1 + f_2}{2} \right\|_{\mathcal{F}} \leq 1 - \delta_{\mathcal{F}}(\|f_1 - f_2\|_{\mathcal{F}}).$$

As a consequence, we have:

$$\frac{f_1 + f_2}{2} - \delta_{\mathcal{F}}(\|f_1 - f_2\|_{\mathcal{F}}) \frac{z_2}{\|z_2\|_{\mathcal{F}}} \in \mathcal{F}.$$

We get:

$$z_2^\top \left( \frac{f_1 + f_2}{2} - \delta_{\mathcal{F}}(\|f_1 - f_2\|_{\mathcal{F}}) \frac{z_2}{\|z_2\|_{\mathcal{F}}} \right) \geq \inf_{f \in \mathcal{F}} z_2^\top f = z_2^\top f_2.$$

Rearranging this last inequality yields:

$$z_2^\top \frac{f_1 - f_2}{2} \geq \frac{\|z_2\|_2^2}{\|z_2\|_{\mathcal{F}}} \delta_{\mathcal{F}}(\|f_1 - f_2\|_{\mathcal{F}}),$$

which implies that:

$$z_2^\top \frac{f_1 - f_2}{2} \geq K \|f_1 - f_2\|_2^q,$$

for some  $K > 0$  independent of  $z_1$  and  $z_2$  since  $\mathcal{Z}$  is compact,  $\|z_2\|_2 \geq A > 0$ ,  $\|\cdot\|_{\mathcal{F}}$  is  $q$ -uniformly convex, and by the equivalence of norms in finite dimensions. By optimality of  $f_1$ , we also have  $z_1^\top \frac{f_2 - f_1}{2} \geq 0$ . Summing up the last two inequalities, we get:

$$(z_2 - z_1)^\top \frac{f_1 - f_2}{2} \geq K \|f_1 - f_2\|_2^q,$$

and (by Cauchy-Schwartz):

$$\|z_2 - z_1\|_2 \geq 2K \|f_1 - f_2\|_2^{q-1},$$

which concludes the proof.